Projects

Table of contents

1. CURRENT PROJECTS

1.1. FACETS: a Framework Application for Core-Edge Transport Simulation

Controlled fusion holds the promise of clean inexpensive energy for the world. I'm working with John Cary's SciDAC II team as a performance engineer to develop a coupled model for fusion reactor simulation. For our first milestones, I'll be partnering with Allen Maloney's TAU developers to characterize the performance of the two key subcomponent codes and to develop an analytical performance model which can be used to drive design decisions as development progresses. More Info.

1.2. Benchmarking and Early Evaluation of High Performance Computing Systems

A portion of my time is invest in vendor interaction, benchmarking, and early system evaluation in support of ORNL's operations, proposal, and procurement activities. You might guess that I've spent far more time puzzling over the ill-behavior of one microprocessor or another than anyone should. ;-)

So you might ask, "Whose computer is fastest?" Don't bother listening to the marketing hype about "peak performance". Comparing the theoretical peak performance of computers is like comparing the tire pressure of automobiles: while low tire pressure is linked to poorer gas mileage, adding air (hot air?) isn't going to make your 1966 Ford Pick-up as efficient as a 2006 Ford Explorer Hybrid. There are simply other factors which are more important.

Similarly with computers, other factors matter more: macro- and micro-architectural features, programming model, software toolchain, and the benchmark itself. Seymour Cray once said that anyone could build a fast processor. The problem is building a fast system. Performance questions boil down to just two fundemental qualities:

- Bandwidth -- How many units of work come out of the system per unit of time?
- Latency -- How long between starting the stream of work and seeing the first result?

You'll notice that I don't reference "storage capacity" or "parallelism" here. Both of these can be viewed as simply tools for managing the latency and bandwidth of the various components of the system. Parallelism allows us to boost bandwidth and amortize latency while storage capacity hierarchies allow the use of lower bandwidth and higher latency components without impacting overall bandwidth.

Thus, rather than just asking about the processor performance, we need to ask about the overall system balance. Does the memory have sufficient bandwidth to feed the processor? Does the processor have sufficient parallelism to hide the memory latency? The same

questions can be applied to the interconnect and the file store. Latency AND Bandwidth.

So it should be a relatively easy task to design a system in which all the latencies are perfectly overlapped with othe work and all the bandwidths are matched through the system from input to output with no bottlenecks... for one application... with no cost constraints. Designing something which has generally well matched bandwidths and well hidden latencies for an arbitrary mix of applications for a cost which is still afforable to the world's largest economic powers is another matter entirely. Good computer architects have my respect.

1.3. Compiler and Tools Evaluations

Compilers (and the rest of the toolchain) are responsible for mapping high level representation of algorithms onto an architecture. As such they are critically interesting to the process of benchmarking and system evaluation. The actual task of analyzing and tuning benchmark codes includes several performance measurement and analysis tasks and some trial and error which lend themselves to some automation.

1.4. CCS Scientific Computing

I'm matrixed to National Center for Computational Sciences Scientific Computing Group, where I share my knowledge and expertise with the users of one of the largest open-science computing facilities in the world.

1.5. Proposal Development

I don't discuss this portion of my work in detail. It isn't classified, but it is business sensitive. However the broad realities of globalization described in Thomas Friedman's best selling book The World Is Flat place the future of our technological leadership (and thus our economic leadership) in a precarious position. Coasting along on the coattails of commodity technology markets cannot assure a leadership position. Someone has to insure our economic and technological future in a world economy rocked by globalization. But who? All this fun isn't free. Send money for research. :-)

1.6. SC07 Industry Exhibits Chair

Chairing the Industry Exhibits for SC07, the joint ACM/IEEE-CS Conference on High Performance Computing, Networking, Storage, and Analytics, to be held in Reno, Nevada.

1.7. SC06 Masterworks Program

Co-Chairing the Masterworks Program for SC06, the joint ACM/IEEE-CS Conference on High Performance Computing, Networking, Storage, and Analytics. The Masterworks program will highlight novel and innovative ways of applying advanced computing, communications and storage technologies to solve challenging, real-world problems. The focus for SC06 will be on solutions that were "beyond our imagination" just a few years ago but are realizable today or in the very near future. As always, Masterworks aims at showcasing applications that have significant impact on the general public. Suggestions for topics and speakers are invited.

2. RECENT PROJECTS

2.1. SC|05 Technical Program

Co-Chaired the Technical Program for the joint ACM/IEEE-CS SC|05 Conference on High Performance Computing, Networking, and Storage. The tech program encompasses many components:

- Invited Speakers. Visionaries and leaders in the field will present their views and prognostications on the field of supercomputing.
- Technical Papers. From hundreds of submissions, a select group of outstanding papers will examine a diverse set of technical issues.
- Tutorials. Many courses, from introductory to advanced level, are offered. Topics taught by renowned experts range across technological trends, specific products, theory and practice.
- Panels and BOFs. Numerous panels and Birds-of-a-Feather sessions will provide forum for discussion, analysis, and a little pontification.
- Workshops. Workshops offer the opportunity for in-depth presentation, discussion and interaction on topics ranging from the highly technical to the geo-political.
- Masterworks. Sessions consisting of invited talks will highlight the applications of supercomputing and associated technologies to real-world problems.
- Awards. Awards ceremony recognizes leadership in supercomputing. Featured are the Seymour Cray Computer Science and Engineering Award and Sidney Fernbach Memorial Award, which include presentations by the winners. Also, the following awards are presented in this ceremony: Gordon Bell, HPC Software Challenge, StorCloud, Bandwidth Challenge, Best Paper, Best Student Paper, and Best Poster awards.
- Scatter/Gather. This year I'm trying something new for SC: a session where we provide open mic time for attendees to describe fresh results, new ideas, and/or works in progress which don't lend themselves to the longer lead time required by other tech program components. It's an experiment!

Thank goodness for co-chairs!

2.2. RAMS Program Intern

The Research Alliance in Math and Science (RAMS) Program is based on the belief that national laboratories and universities, working hand in hand, offer the best opportunity to make a positive impact on the quality of a diverse workforce. The Research Alliance in Math and Science program is designed to provide collaborative research experiences among faculty and students at colleges or universities and DOE national laboratory researchers. These experiences will improve the U.S. competitive research edge while encouraging and promoting Science, Mathematics, Engineering, and Technology (SMET) research throughout the academic year.

2.3. "Lustre™: A How-To Guide for Installing and Configuring Lustre 1.4.1"

Helped organize and lead a five day workshop on Lustre installation and configuration. Using the "redbook" model, a team of Lustre newbies was brought to ORNL from across the US and Europe to write a manual to help those new to Lustre through the growing pains and gotchas of installing and configuring a high performance POSIX compliant cluster filesystem. The document can be found at the LustreUsers.org site along with other guides from our team.

2.4. "Leadership Computing at Oak Ridge National Laboratory"

Co-authored a paper with Studham, White, Fahey, Carter, and Nichols for the 2005 Cray User Group Meeting in Albuquerque, NM. The paper and slides are available at CUG.org.

2.5. Linux Cluster for OS Research

Purchasing the new Linux cluster was easy. So was installing the hardware and software. Shipping and delivery on the other hand entailed a few surprises. The vendor promised the machine in time for the Lustre workshop, but shipped it late enough that it had to be express shipped to arrive on time. However, the one of the crates was *destroyed* prior to delivery at ORNL receiving. Inspection of the damage revealed a four foot by eight foot by three foot crate that had two sides ripped off with the broken pieces piled atop the remains of the crate. The rest of it came apart as soon as the crate was moved to inspect the damage. Amazingly, the equipment inside was undamaged. The hardware and software were installed and the system was brought to demo Linpack within two hours of arrival. A few days later ORNL hosted a very successful Lustre workshop (above) which used the machine as a test platform for lustre installation and configuration. After the workshop, the machine was turned over for

use as an operating system research testbed for the FastOS PetaSSI Project.

2.6. Building 5100 Machine Room

The project to install the Linux cluster started out simply: "put it in the machine room downstairs", they said. What they didn't say was that the room hadn't been used before. Lieberts hadn't been run. No power under the floor. No network. Dirt and metal shavings everywhere. But the ORNL staff is great. A few phone calls mobilized a veritable army of assistance: building stewards, lab space managers, electricians, networkers, and maintenance crews. It all came together fast, but next time someone says to go where no one has gone before, find out if there's power under the floor!

2.7. Analysis of FY06 DOE Budget

If you've every read a Federal Budget document, enough said.